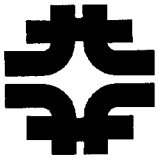

Fermi National Accelerator Laboratory

FERMILAB-Pub-94/232-E

CDF

Performance of the CDF Neural Network Electron Isolation Trigger

B. Denby et. al

Fermi National Accelerator Laboratory
P.O. Box 500, Batavia, Illinois 60510

August 1994

Submitted to *Nuclear Instruments and Methods.*

Operated by Universities Research Association Inc. under Contract No. DE-AC02-76CHO3000 with the United States Department of Energy

Disclaimer

CDF/PUB/CDF/PUBLIC/2755
FERMILAB-PUB-92/232-E

Performance of the CDF Neural Network Electron Isolation Trigger*

B. Denby†, C.S. Lindsey‡
Fermi National Accelerator Laboratory, Batavia, Illinois

M. Dickson§
University of Rochester, Rochester, New York

J. Konigsberg
Harvard University, Cambridge, Mass.

G. Pauletta
University of Udine, Udine, Italy

W. Badgett, K. Burkett, M. Campbell
University of Michigan, Ann Arbor, Michigan

July 23, 1994

Abstract

The performance of the CDF isolated endplug electron trigger in the 1993 run of the CDF experiment is presented. The trigger was designed to select events in $\bar{p}p$ interactions containing isolated electromagnetic clusters in the endplug calorimeter. The trigger was found to be 97% efficient for electrons from W decays in the range of rapidity covered by the trigger, and to provide a background rejection of a factor of 3. The neural network chip used to perform the isolation calculation exhibited no problems with reliability or stability during 8 months of running.

*submitted to *Nuclear Instruments and Methods*
†currently at: INFN Sesione di Pisa, Pisa, Italy
‡currently at: Royal Institute of Technology, Stockholm, Sweden
§currently at: Cornell University, Ithaca, NY

1 Introduction

Three calorimeter triggers, based on the Intel ETANN [1] VLSI neural network chip, were incorporated in the level–2 trigger of the CDF experiment [2] during the 1993 run. The triggers [3] examine calorimeter clusters to see if they are

1. isolated central photons
2. isolated endplug electrons (or photons)
3. semileptonic b decays

These triggers represent the first application of neural network hardware in a major, running high energy physics experiment.

Here we describe the performance of one of the triggers, the endPlug ELEctron isolation (PELE) trigger (see also [4] [5] [6]) which was designed to select $\bar{p}p$ events containing isolated clusters in the endplug [7] electromagnetic calorimeter, coming predominantly from W decays and direct photons.

The cross section for electromagnetic clusters in the CDF calorimeter is about 600 nanobarns for an E_T threshold of 15 GeV, and is dominated by jets with a high fraction of electromagnetic energy; the cross section for real electrons and photons is a tiny fraction of this. In the case of electrons, much of the background can be removed by matching a high p_t track from the fast track finder [8] to the cluster, and this is what is done in the level–2 electron trigger in the central region. This method is not possible in the endplug region because the fast track finder does not extend into this region.

The PELE trigger was designed to reduce the rate from endplug electromagnetic clusters by applying instead an isolation requirement in level–2. Isolation, i.e., the requirement that a cluster have little additional energy surrounding it, is a standard offline variable used in many analyses for electrons and photons. The electrons or photons from W decays and direct photon production will normally not be accompanied by hadrons from fragmentation or gluon radiation, and thus will be well isolated. Since the trigger does not require a track match, it is efficient for photons as well as electrons.

The PELE trigger was run in test mode throughout the 1993 CDF run, with a prescale factor of 12, to check the hardware for accuracy, stability, and reliability. The standard trigger for endplug electrons during the run had a 20 GeV threshold; another trigger with a 15 GeV threshold and a missing E_T requirement was also used. The isolation trigger was formed by

the .OR. of four comparator bits (to be described later), which were read out on every event, allowing the PELE trigger efficiency and rejection factor to be measured even though the trigger was prescaled.

2 Operation of the Trigger

2.1 Trigger Electronics

The neural network triggers were part of an upgrade to the 'conventional' CDF trigger, which has been in operation for several years [2]. In the conventional trigger, DC levels proportional to trigger tower energies (tower size = 0.2 by 0.2 in $\eta - \phi$ space) pass from the calorimeters to receiver boards (called RAW). These levels are then treated by the cluster finder and subsequent elements of the trigger. The ETANN was a natural choice for the neural network upgrade applications because it can use directly the analog signals from RAW, and because it can perform the required pattern recognition in a few microseconds, as required for the level–2 trigger. The neural network hardware is described in greater detail elsewhere [9] [3].

The neural network chips were mounted on ETANN trigger boards [9], which have fastbus control of the chip control voltages, ADC readout of ETANN inputs outputs and control voltages, and DAC's for presenting test patterns to the ETANN. The boards are located in a special fastbus crate, along with a set of shift matrix boards and a control card. On each event, the trigger tower energies from the central and endplug electromagnetic (EM) and hadronic (HAD) calorimeters are carried, as analog levels, on flat cables from the RAW boards to the shift matrix boards. When the level–2 cluster finder passes the address of an EM cluster 'seed tower' to the control board, the shift matrices shift a 5 by 5 tower window centered on the seed tower onto the special 5 by 5 region of the shift matrix array that is hardwired to the crate's analog backplane. The corresponding 5 by 5 HAD window is also shifted. The neural network boards examine the 50 analog signals, corresponding to these two windows, which appear on the analog backplane.

2.2 Isolation Algorithm

An electron striking near the center of a trigger tower will be almost completely contained in the EM compartment of that tower (figure 1a), while

those striking near edges can spread to 2 to 4 towers (figure 1b).[1] Note also that the electrons deposit very little energy in the hadronic compartment of the calorimeter. In contrast, jets distribute their energy over a much larger region of the calorimeter, in both electromagnetic and hadronic compartments. At CDF, jets are clustered with cone sizes of 0.4 or 0.7 in $\eta - \phi$ space, i.e., a diameter of 4-6 trigger towers. It is jets which fluctuate so as to be predominantly electromagnetic which form most of the background to endplug electrons and direct photons.

In *offline* isolation algorithms, a barycenter algorithm is typically used to find the centroid of the cluster, and then the energies of towers within a cone, called the isolation cone, centered on this centroid, are summed. The isolation criterion is then,

$$\frac{E^{TOT}_{0.7(0.4)} - E^{TOT}_{cluster}}{E^{TOT}_{cluster}} < frac$$

where $E^{TOT}_{0.7(0.4)}$ is the total E_T, both HAD and EM, within an isolation cone of radius 0.7 or 0.4, $E^{TOT}_{cluster}$ the HAD+EM energy of the cluster, and $frac$ is the cut on the fraction of cluster E_T permitted outside the cone.

The isolation algorithm executed in level–2 emulates a standard offline isolation cut, but must be able to contend with these complications:

- Only a few microseconds are available to perform the calculation.
- The time constraint requires the η and ϕ of the seed tower to be used as an approximation of the cluster centroid.
- All calculations must be done using trigger tower energies instead of detector tower energies. The trigger towers each contain 6 detector towers.

The solution adopted uses a set of 4 templates defined upon the 5 by 5 trigger tower window as shown in figure 2 (A $5x5$ window is comparable in size to a cone of radius 0.7). The use of templates avoids the need to calculate the isolation cone on each event, which would be time consuming; with parallel hardware, the templates can all be evaluated simultaneously. We define 4 variables, $fisol_i$ by:

[1] The spreading into neighboring towers is more important at the higher η values because the tower size decreases with increasing η.

$$fisol_i = frac * E^{EM}_{i,within\ 2x2} - E^{TOT}_{i,outside\ 2x2}$$

where $i = 1, 4$ is the index of the template. The $fisol_i$ represent the amount by which the energy outside the $2x2$ region exceeds the fraction, $frac$, of the EM energy within the $2x2$ region. A cluster is called isolated in a particular template, i, if $fisol_i > 0$ for that template. This condition is thus similar to the offline isolation criterion. Note, however, that the 'outside 2x2' energy also includes all of the energy in the $5x5$ HAD window, while the 'within 2x2' energy is just EM energy; thus the isolation cut also applies an implicit cut on HAD/EM.

The level–2 isolation criterion is then that $fisol_i$ must be positive in at least one template, i.e.,

$$fisol = max(fisol_1, fisol_2, fisol_3, fisol_4) > 0.$$

The algorithm was optimized on data sets from the 1989 CDF run, including electrons from Z decays and a background sample taken with a low threshold trigger. The optimum value of $frac$ was found in these simulations to be 0.16, which gave 95% efficiency for electrons and a background rejection of a factor of 4. A $1x2$ inner template was also tried but was less efficient because it could not handle showers near tower boundaries (figure 1b). Four templates are necessary since the cluster centroid can often be displaced from the center of the 5 by 5 grid by one tower or more in any direction.

In figure 3 are plotted the values of $fisol$ for a set of endplug W's and for a dataset containing all triggers. The figure shows that the cut, $fisol > 0$ will be very efficient for the $W's$ and reject most of the background. Offline energies were used to estimate the $fisol$ values in the plot, for reasons explained in section 3, where the precise calculation of trigger efficiency and background rejection is done.

2.3 Hardware Execution of the Algorithm

2.3.1 Choice of the ETANN

The Intel ETANN chip was chosen to implement the isolation algorithm in the level–2 trigger. The ETANN is a true parallel analog artificial neural network in which synapse weights are implemented by floating gates [1].

Neural network architectures adapt naturally to template matching applications such as described above, with a processing time which is independent of the number of templates because of the parallel processing.

A very simple network architecture was used, as shown in figure 4. The 50 tower energies form the input layer, and 4 neurons, one per template, are used in the hidden layer. The synaptic weights are set to $frac$ or to -1, so that when the input voltages are applied, a current proportional to the isolation variable, $fisol_i$ appears at the input to the neuron for template i. The network thus is 'hand wired', i.e., no learning algorithm (e.g., backpropagation) was used to determine optimal architecture. This was done intentionally (in spite of the potential for improved performance using backpropagation) since it was desired to have a network that was readily understandable in terms of simple cuts.

For the PELE trigger, the ETANN chip was run in HGAIN mode, in which the neuron transfer function is essentially a step function[2], so that when the isolation criterion for a given template was satisfied, the corresponding neuron gave a 'high' value at its output. The neuron outputs are followed on the ETANN board by comparators; in HGAIN mode these are not really needed and are simply used for level conversion. The final trigger bit is an .OR. of the 4 template responses.

2.3.2 Alternatives Considered

It would be straightforward to implement the PELE trigger circuit using simple comparators instead of neurons, and resistors instead of floating gate synapses. In fact, a board with such a discrete construction was built and found to function well in benchtop tests. The ETANN solution was retained, however, because the chip has programmable weights, which allow the value of $frac$ to be modified via computer, and because of the flexibility to extend to more complicated template schemes should the need arise. Also, the ETANN boards have many desirable features such as computer setting and monitoring of control voltages, on board DAC's for sending test patterns to the trigger, etc.

[2] In many neural network applications, including the photon isolation and b triggers, the neurons have transfer functions with a smooth, sigmoidal shape.

2.3.3 Implementation Challenges

Although the network is hand wired, loading the algorithm into the chip presented a number of challenges (see also [6]):

- The synapse multiply characteristic is not perfectly linear [1].
- The weights are limited to the range ± 2.5 and have a precision of about 6 bits.
- Tower energies from 0 to 40 GeV must be scaled to lie within the ETANN input range of 0 to 3 volts without loss of accuracy.

Another problem stemmed from the fact that, in most applications, the ETANN is run in bipolar mode. That is, a control voltage, V_{refi}, with a value of one half of the voltage swing of the inputs, is subtracted from each input voltage before it reaches the synapses, with the anticipation that most of the inputs will normally run at about mid–range. In the case of electron clusters, especially isolated ones, most of the 50 energies are in fact very small, which implies a large negative voltage into the synapses once V_{refi} is subtracted. The large quiescent input voltage implies a large quiescent current in the neuron amplifiers, which generates a substantial amount of heat. To reduce the heating, V_{refi} was set to the relatively low value of 0.2 volts. It was not possible to set it exactly to zero because this produces nonlinear behaviour. In this situation, the quiescent level of the inputs with zero energy is -0.2 volts.

When summed over the 50 towers, these offsets produce a substantial shift in the $fisol_i$ current at the templates. This shift was compensated by introducing individual biases for each of the neurons. The biasing was done using the 14 unused ETANN inputs, connected to the neurons with weights set near the maximum value of 2.5, to which programmable DACs were connected. It should be noted that this biasing of the neurons is not just a characteristic of the ETANN; it was also found necessary on the discrete board mentioned above to remove offsets in the comparator thresholds.

Test patterns of isolated and non–isolated clusters, generated with a simple Monte Carlo, were sent to the ETANN through fastbus and the values

$$frac_i = \frac{E^{TOT}_{i,outside\ 2x2}}{E^{EM}_{i,within\ 2x2}}$$

were calculated for each template. When $frac$ is less than 0.16, if everything is working properly, the cluster is isolated and the output should be high; when greater than 0.16 the output should be low. A scatter plot of output versus $frac_i$ for a pattern set is shown in figure 5. The bias DAC values were adjusted until the transition from off to on of the neurons occurred as closely as possible to $frac_i = 0.16$. The plots show a finite width of the transition region, rather than a sharp edge, due to noise, finite resolution of the analog electronics, non–linear synapse multiply characteristic, etc.

The biases had been initially determined in the Intel training system [10]. When the chip was transferred to the ETANN trigger board it was found necessary to readjust slightly these biases, and during collider running, the biases had to be readjusted again[3]. The width of the transition region was also somewhat larger during collider running; this effect will be discussed again in the section on background rejection. The reason for the differences between collider–on and collider–off modes has not been understood; however, once the correct biases for collider running were established, they remained very stable.

3 Performance of the Trigger on Endplug W's in the 1993 CDF Run

3.1 Data Set

A sample of 3.5 pb^{-1} of data taken in January and February of 1993 was used to measure the efficiency of the PELE trigger for W's in the endplug. Because the comparator bits were read out on every event, events for which the isolation trigger bit itself was not acknowledged could be used in this study. The standard analysis cuts for the endplug W sample were applied to select a clean sample of W's. Those cuts are:

1. $E_T^{cluster} \geq 20\ GeV$
2. $E_T^{miss} \geq 20\ GeV$
3. $1.1 \leq |\eta_{cluster}| \leq 2.2$
4. HAD/EM of cluster ≤ 0.05

[3]This was true in spite of having carefully maintained power supply voltages identical in both environments. Failure to do so is a known way of producing offsets on the ETANN.

5. track match in forward tracking chamber

6. pad $\chi^2 \leq 3$.

7. energy outside cone of 0.4 $\leq$ 10% of $E_{cluster}^{TOT}$

8. $40 \leq M_T^W \leq 100\ GeV/c^2$

where $E_T^{cluster}$ is the transverse energy of the cluster, E_T^{miss} is the missing transverse energy in the event, and M_T^W is the transverse mass of the W. The pad χ^2 is a measure of how well the transverse profile of the cluster matches the mean electron profile. The seventh cut is the offline isolation cut. These cuts are on *offline* E_T, which can differ from what is actually seen by the level–2 trigger.

3.2 An Endplug W Found only by PELE

Figure 6 shows the transverse and longitudinal calorimeter energy distributions for an event satisfying the above W cuts and passing *only* the PELE trigger, even though that trigger was prescaled by a factor of 12. The event was missed by the 20 GeV trigger and 15 GeV plus E_T^{miss} trigger due to E_T and/or E_T^{miss} threshold inefficiencies in those triggers. One of the original purposes of the PELE trigger was in fact to help correct for the turn-on of the 20 GeV threshold trigger. The event contains an electron with an offline E_T of 26 GeV at an η of -1.3, matched by a track in the central tracking chamber, E_T^{miss} of 22 GeV, and M_T^W of 52 GeV, and is consistent with being a low E_T endplug W.

3.3 Trigger Efficiency

The PELE trigger flags clusters which are isolated according to the level–2 algorithm by turning on one or more of the comparator bits corresponding to the 4 isolation templates. The trigger efficiency is defined as the number of W candidates that have at least one comparator bit on divided by the total number of W candidates. With this definition, the PELE trigger is found to be 97% efficient for W's.

When seed towers are near the minimum and maximum η values of the endplug, parts of the level-2 isolation templates will extend into the central or forward calorimeters respectively. Since the forward calorimeter is not implemented in the neural network triggers, it is not possible to use the isolation templates when the seed tower is less than two towers from the

edge of the plug (i.e.,$|\eta| > 2.|$). Since it was desired to measure the efficiency of the neural network hardware, without including geometric effects, these latter cases were excluded from the efficiency study. This eliminates about 19% of endplug W's, however many of these are not usable at CDF anyway because the tracking efficiency falls rapidly for $|\eta| > 1.7$, precluding accurate momentum and charge meaurements.

Because of the 20 GeV offline threshold, it was not possible to measure directly the PELE trigger efficiency below 20 GeV; however, any differences in efficiency from that above 20 GeV should be due only to non-linearities in the ETANN synapse multiply characteristic. Studies using the Intel trainer showed no change in trigger performance for cluster energies down to 10 GeV. The efficiency that we have measured is thus probably also valid for the 15 GeV threshold at which the trigger was intended to run.

3.4 Background Rejection

It was not possible to make a scatter plot like figure 4 using the W sample because the backplane voltages were stripped off from this sample in the offline processing. We can, however, get an idea of the sharpness of the trigger turn-on in the $fisol$ variable for collider running by looking at some background data for which the backplane values were saved. Figure 7 shows the turn-on of the trigger efficiency in the $fisol$ variable for this data, using the energies read from the analog backplane to calculate $fisol$. To make the plot, the ratio of two histograms of number of events versus $fisol$, for trigger *on* and for all events, was taken. Note that that the turn-on in efficiency is not perfectly sharp: 50% efficiency is attained at about $fisol = -2.5\ GeV$, full efficiency near $fisol = 0$. This is due to noise and the analog nature of the calculation as discussed above.

The non-perfect turn-on of the trigger is undesirable because it causes it to pass more background, as background clusters normally have $fisol$ negative. The simulations of the trigger using 1988/89 data had predicted a cross section of 150 nb, i.e., a reduction of the background of a factor of 4. During actual running in 1993, the unprescaled rate of the PELE trigger was 200 nb, corresponding to a reduction factor of 3, due to the non-perfect turnon. It is possible to lower this rate to 150 nb by raising the neuron biases a few millivolts, effectively shifting the turnon in $fisol$ slightly in the positive direction, degrading only slightly the efficiency for W's, to about 93%.

4 Stability and Reliability

During collider running, the neural network triggers were monitored with the standard online trigger monitoring program. Histograms of the number of firings of each template were compared with a standard histogram. Although there were problems from time to time with various parts of the system, which introduced anomalies in these histograms, the neural network chips themselves were never found to be at fault.

The rates of all triggers were monitored during running with an online program which gave an alarm when rates went outside tolerances of 10%. The PELE trigger rate is very sensitive to the *fisol* threshold due to cross section for EM clusters which rises very sharply with increasing E_T. Bench-top studies before the run had indicated a temperature dependence of the *fisol* threshold. During running, the ETANN boards were in a standard fastbus cooling system, and no problems with stability of the trigger rate were observed during the 8 months of running.

4.1 Future Use of Neural Net Triggers on CDF

In spite of its good performance, the PELE trigger will not be used in future runs. The plug electron physics group, which had initially requested the trigger, opted instead to use the trigger with the E_T^{miss} requirement as the principal endplug electron trigger, because it was feared that an isolation cut in the trigger would bias the W mass measurement. The central photon isolation trigger worked very well [11] and will be retained for future running. A new trigger to detect τ particles, using the same hardware, is also being developed [12].

5 Conclusion

The CDF neural network isolated endplug electron trigger was found to be 97% efficient for good W's within the η coverage of the trigger. It was able to trigger upon W's which were missed by the other plug electron triggers due to their threshold inefficiencies. The trigger has some residual sensitivity for non-isolated events, which degrades the predicted background rejection factor of 4 to the observed factor of 3. A factor of 4 rejection can be achieved with a small change to the neuron biases, with only a small decrease of efficiency. Thus the chip performance is quite close to that predicted in

the simulations. The neural network chips showed no failures or stability problems during 8 months of running at the Fermilab Collider.

References

[1] Data Booklet for Intel 80170NX *Electrically Trainable Artificial Neural Network*, Intel Corp., 2250 Mission College Boulevard, MS SC9-40, Santa Clara, CA, June, 1991.

[2] M. Campbell et al., *Nucl. Instr. Meth.* A265 (1988) 326-335.

[3] W. Badgett et al., *A Neural Network Trigger Used in CDF*, Proceedings of IEEE Nuclear Science Symposium, October, 1992, Orlando, FL.

[4] B. Denby et al., *Proposal for a Level-2 Isolated Plug Electron Trigger for the 1991/2 Collider Run*, CDF Experiment Technical Note CDF/DOC/CDF/PUBLIC/1538, unpublished.

[5] B. Denby, *Status of Neural Net Triggers at the Fermilab Tevatron*, in *New Computing Techniques in Physics*, D. Perret-Gallix, Ed., World Scientific, 1992.

[6] C.S. Lindsey and B. Denby, *A Study of the Intel ETANN VLSI Neural Network for an Electron Isolation Trigger*, CDF Experiment Technical Note CDF/DOC/CDF/PUBLIC/1850, unpublished.

[7] Y. Fukui et al., *Nucl.Instr.Meth.* A267:280,1988.

[8] G.W. Foster, J. Freeman, C. Newman-Holmes, J. Patrick, *Nucl. Instr. Meth.* A269:93,1988.

[9] W. Badgett, M. Campbell, D.Wu, *Neural Net Based Calorimeter Trigger*, University of Michigan Trigger Group, unpublished.

[10] iNNTS Intel Neural Network Training System User's Manual, Intel Corp., Santa Clara, CA.

[11] W. Blair, private communication.

[12] J. Conway, private communication.

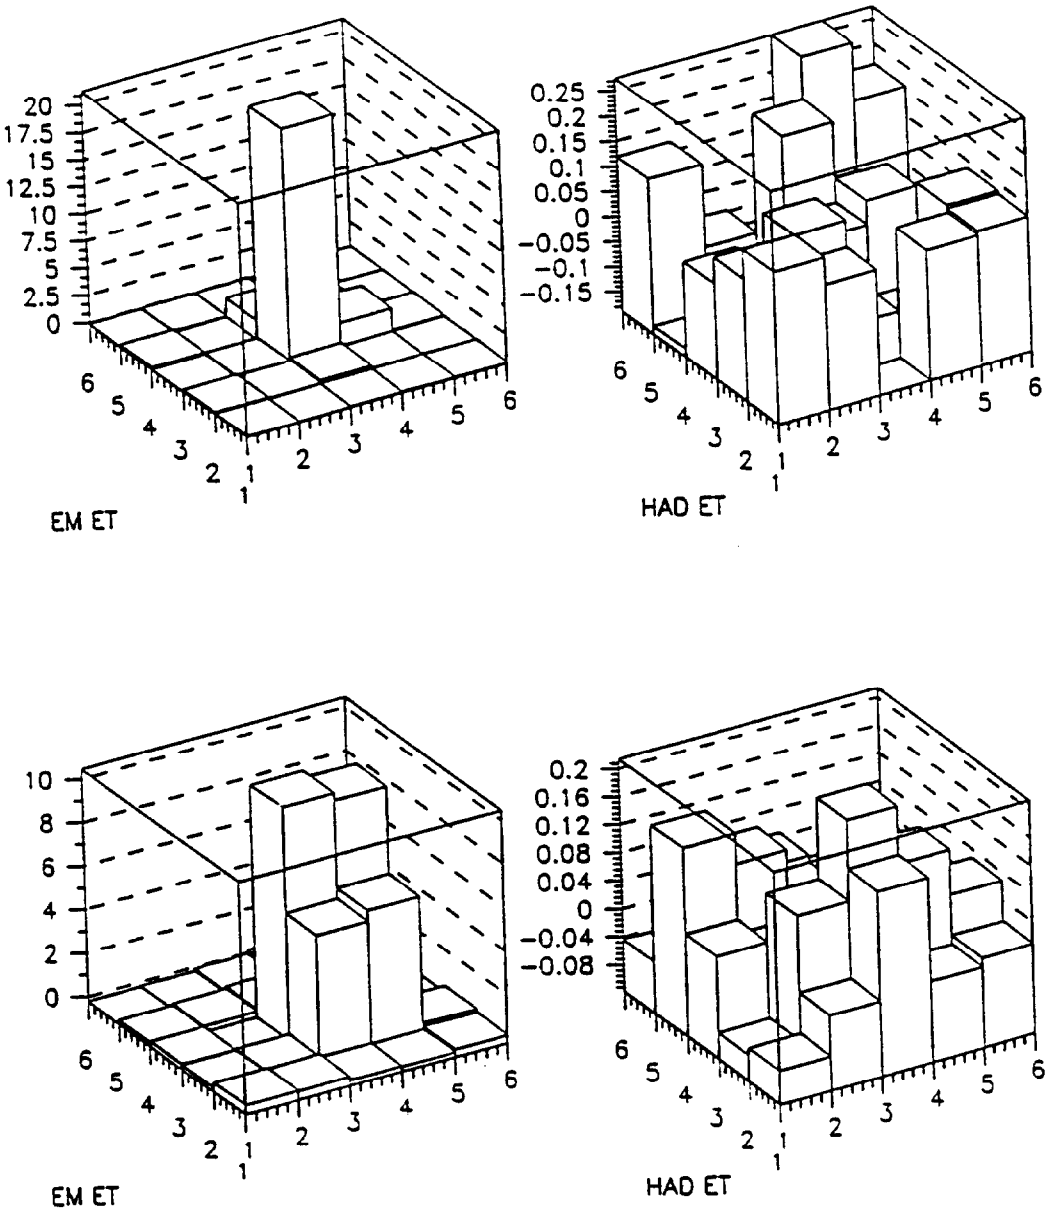

Figure 1. Two types of electron shower profiles in the endplug calorimeter, from actual CDF data. In a), the electron strikes near the center of the seed tower and deposits most of its energy in the electromagnetic (EM) compartment of that tower. In b), the electron strikes at a corner, and shares its energy among 4 towers. In both cases, very little energy is deposited in the hadronic compartment of the calorimeter (note vertical scales).

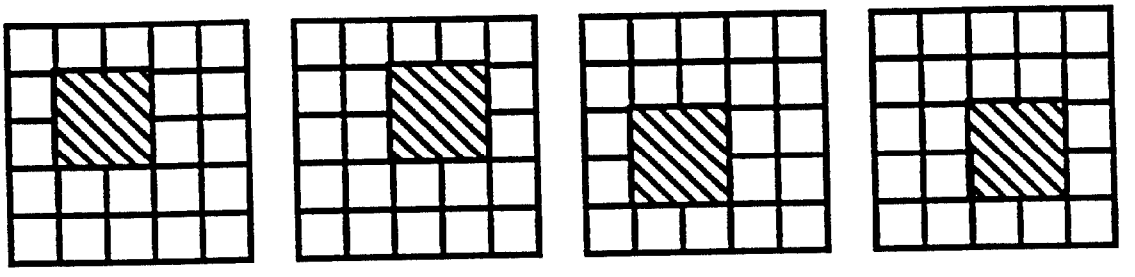

Figure 2. The four isolation templates. The grid represents a 5 by 5 region of the EM calorimeter centered upon the seed tower. In each template, white square energies are weighted with the value -1, blacks squares with *frac*. The cells in the HAD compartment (not shown) are all weighted -1. With these weights, the isolation variable *fisol* (described in text) will appear at the input of the neuron corresponding to the template.

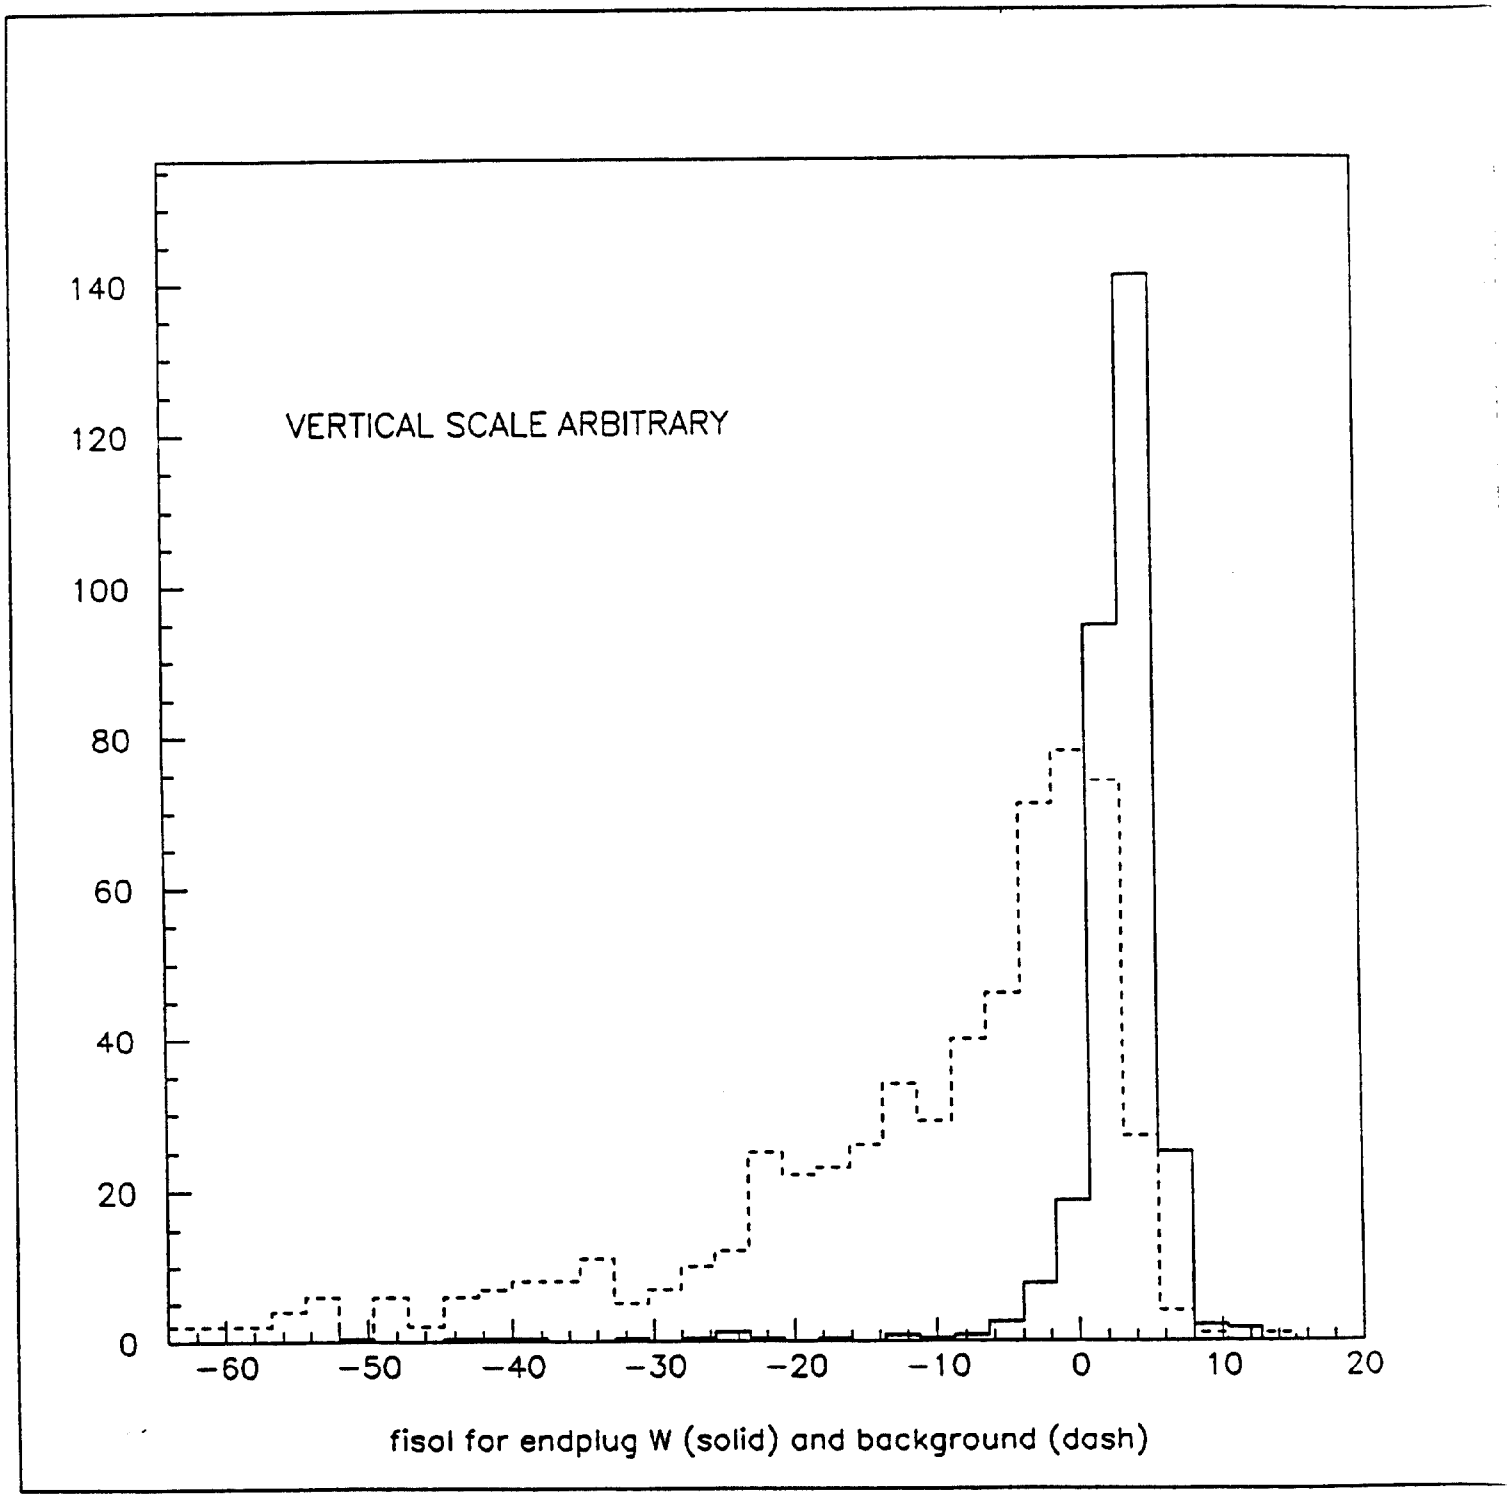

Figure 3. ***fisol*** **values (calculated from offline energies) for a set of endplug** W**'s (solid curve; cuts are explained in section 3) and for a dataset containing all triggers (dashed curve). The vertical scale is arbitrary. The PELE trigger is designed to select events with** $fisol > 0$**.**

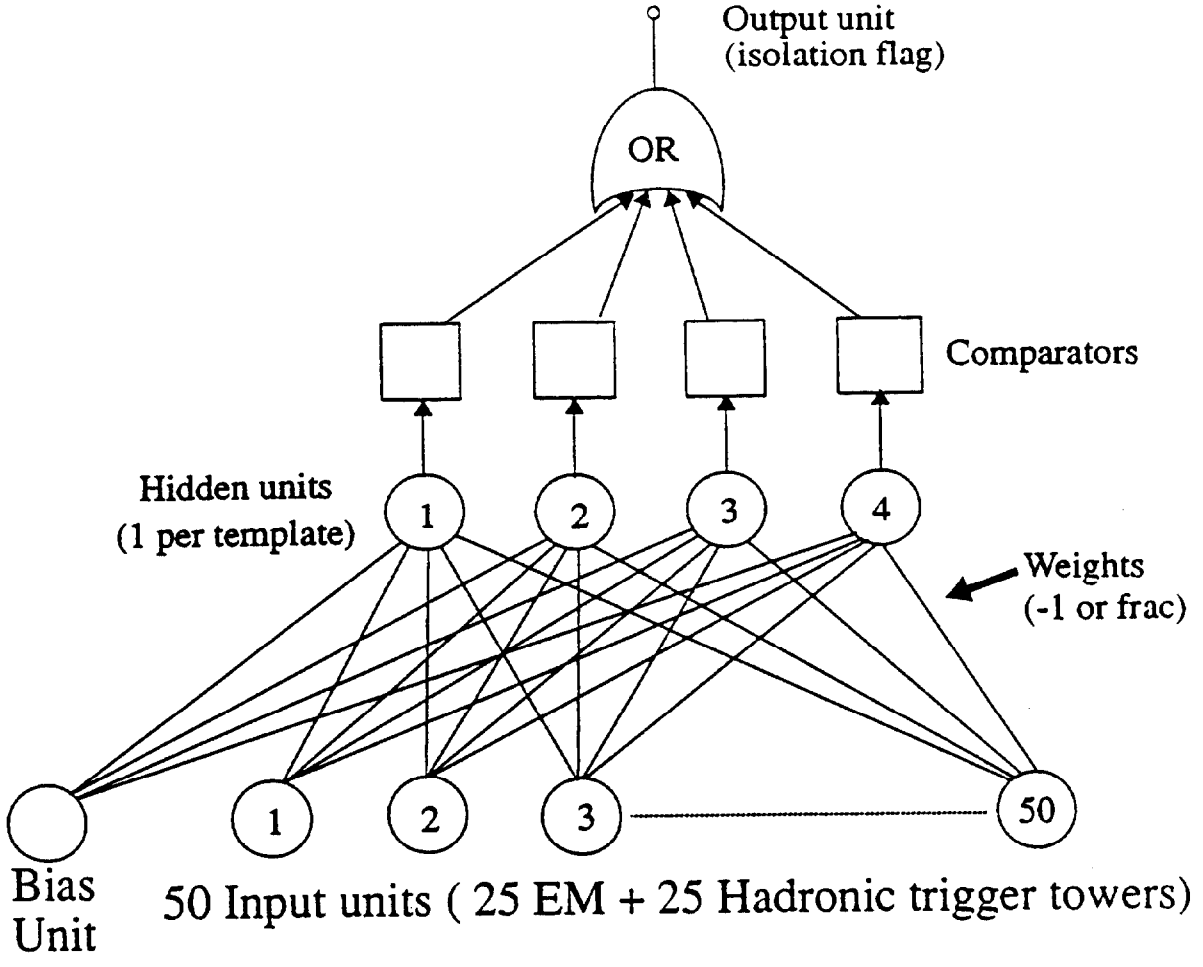

Figure 4. Schematic representation of the feed forward neural network architecture that calculates the isolation algorithm.

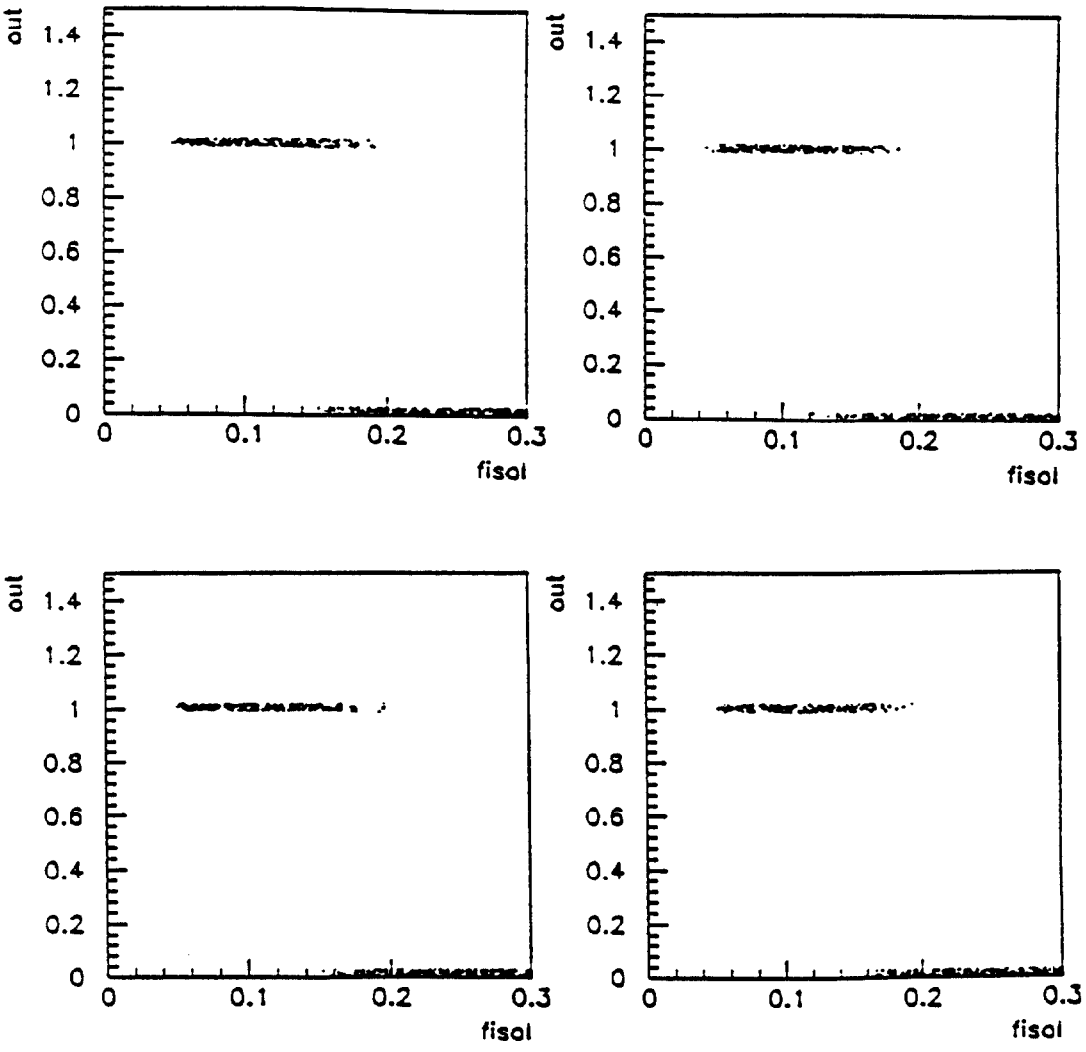

Figure 5. Scatter plot of neuron output versus *fisol* for each of the 4 isolation templates as measured with the ETANN in the Intel trainer, for a set of Monte Carlo endplug clusters. The transition from *on* to *off* occurs at *fisol* = 0.16, for each template, as it should; however, a finite width of the transition region is visible, due to imperfections in the analog computation.

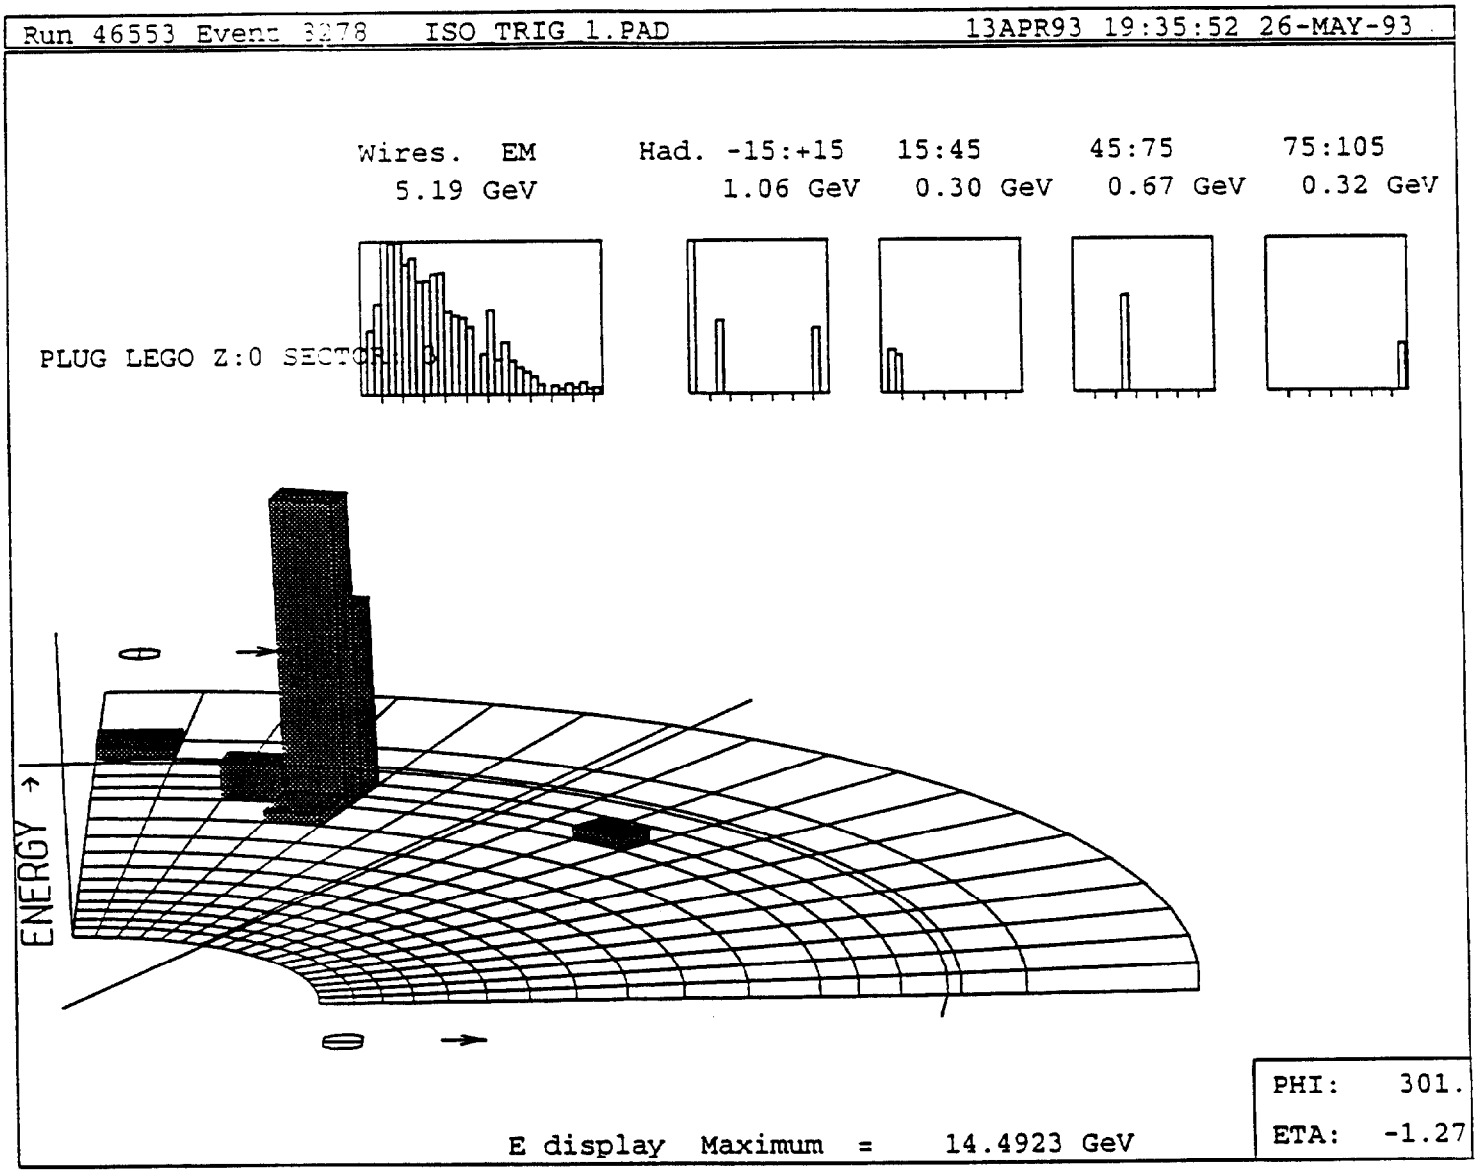

Figure 6. An endplug W event from the 1993 CDF run, detected by the PELE trigger but missed by the standard plug triggers. The lower part of the figure is an $r - \phi$ lego plot of the energy deposit of the electron in the endplug calorimeter. In the boxes above is shown the longitudinal profile of the electron energy deposit.

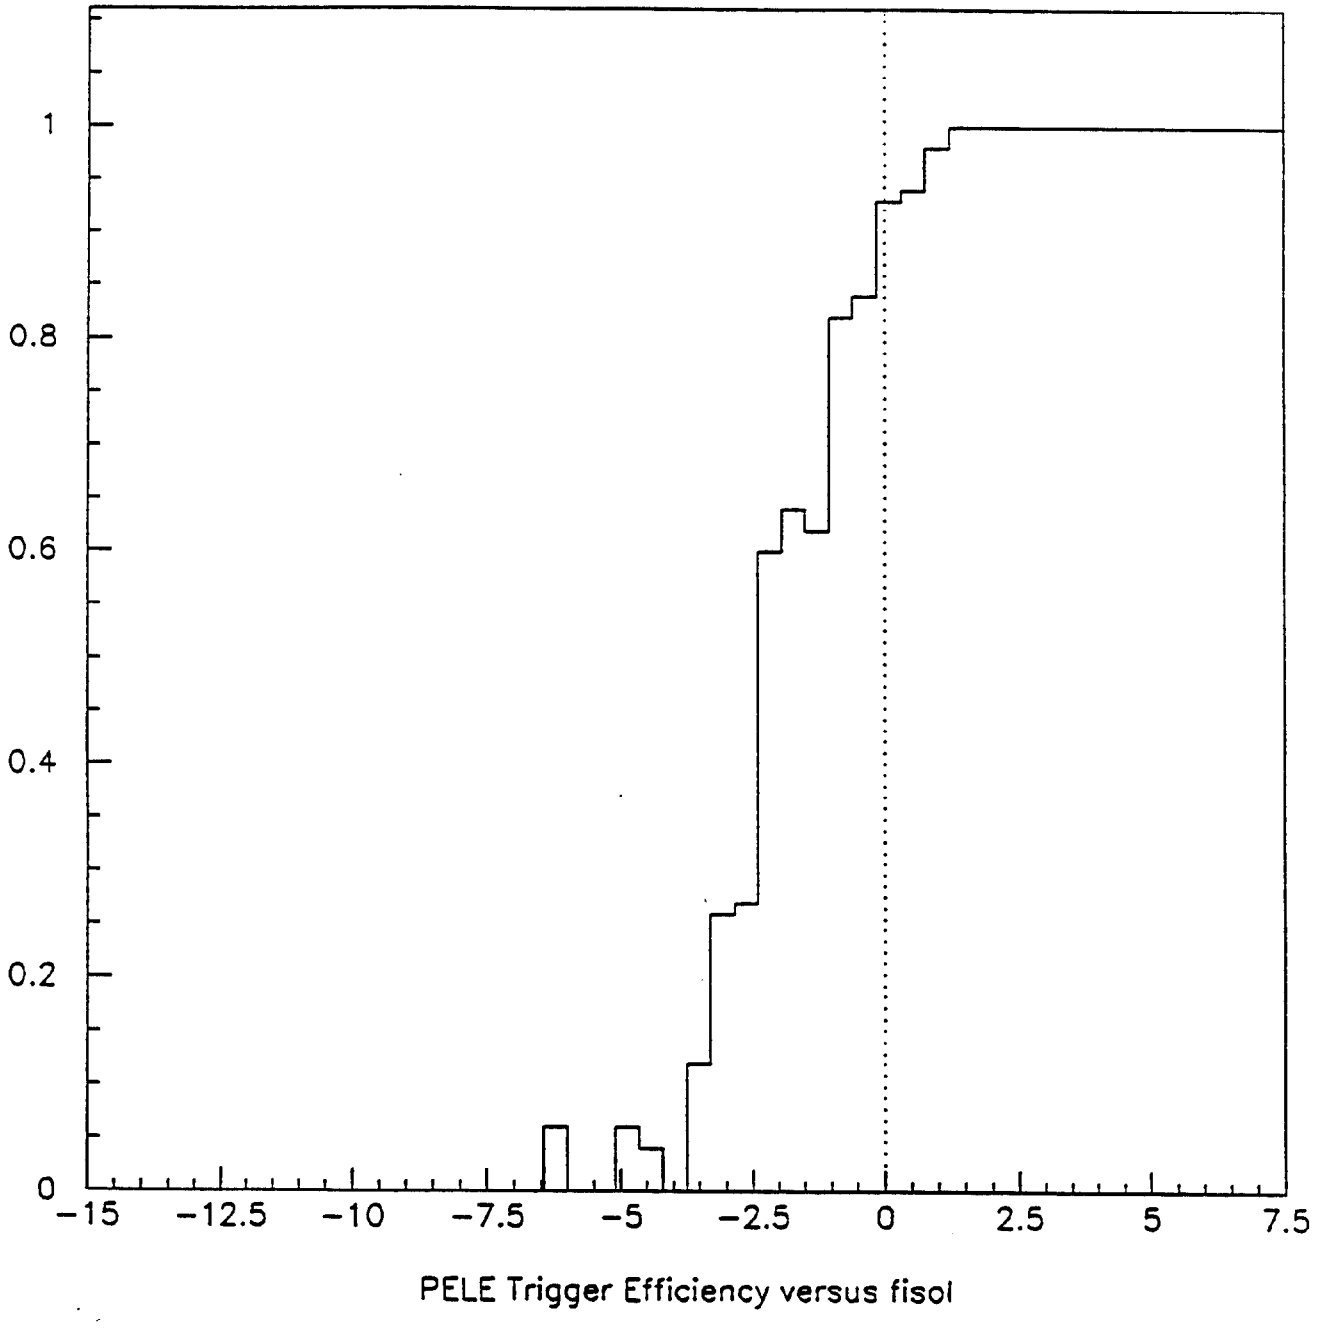

Figure 7. Efficiency of the PELE trigger as a function of the variable *fisol*. With perfect electronics this plot should be a step function at zero. The residual sensitivity for *fisol* < 0 degrades the background rejection factor, as discussed in the text.